INTEGRATED COLOR CODING AND MONOCHROME MULTI-SPECTRAL FUSION

Tamar Peli, Ken Ellis, Robert Stahl[*]
Atlantic Aerospace Electronics Corporation
470 Totten Pond Road, Waltham, MA 02154

Eli Peli
Schepens Eye Research Institute, Harvard Medical School
20 Staniford Street, Boston, MA 02114

ABSTRACT

This paper describes a new integrated color coding and a contrast-based monochromatic fusion process. The fusion process is aimed for on board real time application and it is based on practical and computationally efficient image processing components. We developed two methods for color coding that utilize the monochrome fused image. Each of the color coding methods provides consistency of color presentation as a function of daytime, background variability and illumination conditions. The new monochrome fusion process maximizes the information content in the combined image, while retaining visual clues that are essential for navigation/piloting tasks. The method is a multi scale fusion process that provides a combination of pixel selection from a single image and a weighting of the two/multiple images. The spectral region is divided into spatial sub bands of different scales, and within each scale a combination rule for the corresponding pixels taken from the two components is applied. Even when the combination rule is a binary selection the combined fused image may have a combination of pixel values taken from the two components at various scales since it is taken at each scale. We also applied a combination rule that takes a weighted sum of the two pixel values.

The fusion concept was demonstrated against imagery from image intensifiers and forward looking infrared sensors currently used by the U.S. Navy for navigation and targeting. The approach is easily extendible to more than two bands. To be effective, the fused imagery maintains relationships that correspond to natural (daytime) vision of the same features. Under stress the human operator is liable to revert to his most natural interpretation and act on it. Thus any image transformation that distorts such relationships, even if it can be learned by the user and responded correctly in the lab or during training, may be less effective. The fusion process provides substantially different colored fused image, which is better tuned to the natural and intuitive human perception. These are necessary for pilotage and navigation under stressful conditions, while maintaining or enhancing the targeting detection and recognition performance of proven display fusion methodologies.

1.0 INTRODUCTION

Significant improvements were made in night vision devices/sensors over the last two decades to aid military forces in conducting night operations. There are two types of night vision devices currently in use; image intensifiers (I^2) and forward looking infrared (FLIR) sensors. Each sensor provides a monochrome image that can be used for either navigation and/or targeting. Image intensifiers collect reflected energy in the 600 to 900 nm range, while FLIRs collect emitted energy from the far infrared spectrum (typically 8 to 12?). Targets as well as terrain features may exhibit noticeable contrast in either domain dependent on a variety of conditions and environments. The two sensors tend to excel under different conditions and as a result of their complementary nature there may be a strong desire and benefit from having imagery from the two sensors during a night flight mission.

---

*former employee

In addition to improved visibility at night, detection performance of targets in concealment camouflage and deception (CC&D) conditions is greatly improved, if multi-spectral imagery is available. This was demonstrated

Form SF298 Citation Data

Report Date ("DD MON YYYY") 00001999	Report Type N/A	Dates Covered (from... to) ("DD MON YYYY")

Title and Subtitle Integrated Color Coding and Monochrome Multi-Spectral Fusion	**Contract or Grant Number**
	Program Element Number
Authors Peli, Tamar; Ellis, Ken; Stahl, Robert; Peli, Eli	**Project Number**
	Task Number
	Work Unit Number
Performing Organization Name(s) and Address(es) Atlantic Aerospace Electronics Corporation 470 Totten Pond Road Waltham, MA 02154	**Performing Organization Number(s)**
Sponsoring/Monitoring Agency Name(s) and Address(es)	**Monitoring Agency Acronym**
	Monitoring Agency Report Number(s)
Distribution/Availability Statement Approved for public release, distribution unlimited	
Supplementary Notes	
Abstract	
Subject Terms	
Document Classification unclassified	**Classification of SF298** unclassified
Classification of Abstract unclassified	**Limitation of Abstract** unlimited
Number of Pages 9	

[Peli 1997a] for automatic target detection using multi-spectral imagery collected by the ERIM M7 and Daedalus 1268. For a human operator the multiple sources of imagery need to be fused and displayed in a form that is easy and natural to interpret and that will result in improved targeting and navigation performance.

Our long-term objective is to develop a robust color coding method that utilizes the output of a monochrome multi-spectral fusion process. This paper describes the new contrast-based monochromatic fusion process that maximizes the information content in the combined image while retaining visual clues that are essential for navigation/piloting tasks. The monochrome fusion process selects features from each band based on visually relevant contrast measures, but performs the fusion in the amplitude domain. The resulting algorithm was tested by fusion of imagery from image intensifiers and forward looking

Infrared sensors currently used by the U.S. Navy for navigation and targeting. The monochrome fusion stage was tested against a wide range of imagery conditions, and was shown to capture the details from both inputs using a single parameter setting (no image-by-image tuning). The fused imagery was also shown to preserve the depth (shading) information of the visible input. We developed multiple methods for color coding which are derivatives of our technique of polarity preserving color coding. These methods will be reported separately.

1.1. Current Fusion Technologies

Current fusion procedures can be divided into multiple categories based on two main variables; color vs. monochrome, and single scale vs. multi-scale. The use of color in image fusion was frequently advocated under the argument that color contrast can provide improved detection performance when added to luminance contrast. This was demonstrated for very low, near threshold contrast [Gur 1993]. This improvement, however, disappeared when contrast is increased. In a targeting function, both detection and discrimination are required. The contrast level required for discrimination may be sufficiently higher to reduce the noticed improvement in detection performance. The investigation of Perconti et al [Perconti] also indicates that the benefits of color representation of fused imagery is not clear for pilotage tasks, but may be of utility for targeting.

A typical color fusion approach presents the multiple (monochrome) bands or a derivative of them using color displays by transforming (projecting) them onto display variables such as the RGB or the Luminance-chroma-saturation. This approach is taking advantage of the observer's color vision to introduce additional dimensionality for interpretation. For up to three bands, the apparent advantage is that one does not have to make a choice and discard information - all of the information from the multiple bands can be presented in a way that the visual system will make the selection/decision itself. Both the Naval Research Lab [Scribner 1993] and Lincoln Laboratory [Waxman 1995] have developed color fusion procedures for human observers.

The monochrome approach generates a single image from the multiple (typically two) bands. The rules for combination vary from selecting pixels from a single image to weighting of the two images [Morgan 1991]. The combination rule/decision may be based on the amplitude to make this selection (i.e. Paval et al), or on a contrast measure [Toet 1992], or on a contrast-like measure [Waxman 1995]. In either of these cases the fusion is performed in the decision domain (e.g. fusion in contrast space if decision is based on contrast).

Once a mode of display (color or monochrome) has been selected, the decision is between "single" scale and multi-scale image decomposition for fusion. Waxman's [Waxman 1995] shunting approach is essentially a bandpass filtering method. Each band is processed with such filtering to yield a single scale, which is combined using the same method with one of the other bands. The specific spatial scale to be selected is not specified in the various reports, but it is determined by a few of the free parameters defined for this method. The application of the Peli and Lim algorithm [Peli 1979] for multi-band fusion developed by [Scrofani 1997] also represents a single scale (possibly wide band) fusion of the two images. These single scale approaches can be extended to a multi-scale fusion variant.

The single and multiple scale fusion methods discussed above include explicitly a spatial preprocessing or "enhancement" stage before fusion [Toet 1992; Waxman 1995; Ryan 1995]. The preprocessing may be needed for two reasons. Since the two sensors may have widely varying dynamic range and result in highly varying image histograms and contrast, some normalization is required before the bands can be fused into a common space. We

will refer to this process as "normalization". In addition, the spatial preprocessing itself may be used to improve the visibility of desirable features in the image and thus serves as an image enhancement module. Waxman's processing is based on the model of early processing stages of the visual system. Toet [Toet 1992] achieved normalization by combining "contrast" rather than amplitude, which was shown by Peli [Peli 1990] to provide a form of enhancement. In particular, this tends to increase the visibility of low-contrast features in low-luminance areas of the image (in shadows). At the same time, as shown by Peli [Peli 1990], this approach reduces the visibility of details in high luminance areas. The latter problem can be resolved with the proper set of parameters applied to the Peli and Lim adaptive enhancement algorithm as was demonstrated in [Ryan 1995] at the pre-fusion processing stage. Note however that the Peli and Toet enhancements are parameters free, while the Peli-Lim and Scrofani approaches require specific tuning of the enhancement parameters.

The method of multi scale fusion by [Paval 1991] provides a combination of pixel selection from a single image and a weighting of the two/multiple images. In Paval's approach, the spectral region is divided into spatial sub bands of different scales, and within each scale a combination rule is applied for the corresponding pixels taken from the two components. Even when the combination rule is a binary selection [Toet 1992], the combined fused image may have a combination of pixel values taken from the two components at various scales since it is taken at each scale. A combination rule that takes a weighted sum of the two pixel values has also been applied [Paval 1991].

Perconti et al [Perconti] have recently reported the result of task evaluation of various methods of two band fusion. They found that object recognition tasks (targeting) are performed best with the FLIR alone and with the Waxman color fusion method that incorporated the FLIR image, however the value of the color aspect was not clear. In evaluating the horizon perception task, no difference was found between the various imaging and presentation modes. In the geometric perspective task, which is important in navigation, the monochrome version of Waxman provided faster response. Scribner's color format that maintains the phase (polarity) for the visible image and separates it from the IR by color was found significantly better than Waxman color algorithm and not different than the Waxman monochrome version. None of the formats were found to be desirable for night helicopter pilotage from the short video segments. They concluded also that the color fusion might have its greatest utility as a targeting aid.

Ryan et al [Ryan 1995] reported on a comparison of FLIR only and Texas Instruments proprietary monochrome fusion algorithm (FLIR & intensified). In extensive field tests they found a substantial preference for the fused configuration for the specific pilotage maneuvers and for overall ranking. It should be noted that the study found no significant differences in any of the tasks and measures between the Lincoln Lab color and monochrome versions of the algorithm.

1.2   Multi-Spectral Fusion Process

Our fusion process for two bands is a two step process; monochrome fusion followed by pseudo chromatic color enhancement. The overall approach is illustrated in Figure 1.1.
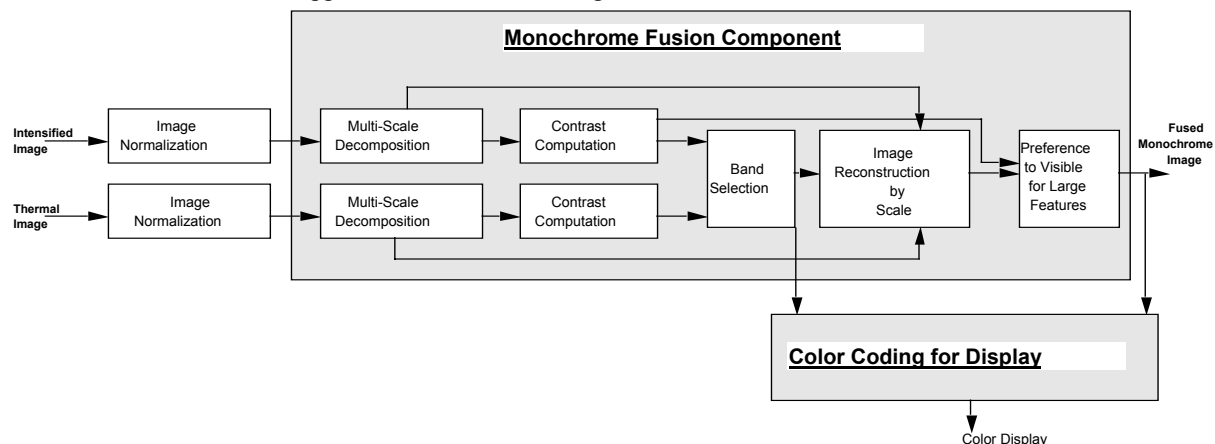

Figure 1.1: Fusion Process for Two Bands

The monochrome fusion has three key attributes

    1) scale-by-scale fusion using oriented filters

2) decision based on contrast in each scale, fusion in amplitude domain
3) preference for visible band at least for large scales to preserve shape from shading information.

Points 2 and 3 above distinguish our process from others in the literature. The color-coding illustrated in Fig 1.1.is applied to the monochrome fused image. The color-coding aims to help the user identify the identity of the band that contributed to colored features.

2.0 METHODS

2.1   Imagery for Algorithm Development and Evaluation

Two sources of imagery were used in the study; imagery provided by the Naval Post Graduate School (NPS) and existing collections of static multi-spectral imagery (collected by the Daedalus 1268 MS sensor). The sets of images obtained from NPS included static images and video segments collected by an early prototype fusion sensor system developed by Texas Instruments and the Night Vision Electronic Sensor Directorate.

2.2   Image Normalization

The fusion process is applied to normalized imagery. The two bands may result in images with differing characteristics including dynamic range and contrast range. The goal of the normalization stage is to bring the multiple outputs from the different bands into a common framework. Since an enhancement may be simultaneously achieved, depending on the normalization scheme, the normalized/enhanced imagery has been used for fusion evaluation to assure critical evaluation of the fusion results separately from any other processing effect. Thus, the fused image was compared to the individual normalized inputs and not to the raw imagery.

It is possible to apply either global or local image normalization procedures. We applied a global histogram based normalization procedure. The normalization of images was accomplished by computing a global mean (m) and standard deviation (?) for each image. The statistics ignore a percentage of the two tails of the histogram. Once these statistics are obtained, each pixel (i,j) is normalized using the following equation:

$$p_o(i,j) = \frac{D}{2} + \mid \frac{2^n}{2n_o} \mid \lceil \frac{p_i(i,j) - m}{\sigma}$$

where p_i is the input, p_o the output, D the maximum range of the data (for eight bits this would be 256), n the desired bits of dynamic range and the number of standard deviations to cover the dynamic range.

2.3   Image Registration

The video images suffered from considerable misregistration between the two sources. There was a mismatch in both aspect and field of view. For the static images, registration mismatch was also observed; however, this was only in the form of translation. As might be expected, the accuracy of registration is key to successful image fusion, misalignment reduces the sharpness and contrast of the fused images. Using our alignment process described below, we were able to register these image pairs and achieved a substantial improvement in fused image quality.

Since the relative geometry between the two sensors was unknown, we warped the images using manual selection of tie-points. Common points were selected in both images and these coordinates were used to warp one image to the other (visible to IR). Selecting tie-points can be made difficult due to the mismatch in the image content and contrast differences, therefore an automatic scheme has been implemented to fine-register the images after the manual warping. To register the motion video imagery, a single set of manual tie-points was chosen and used to warp the visible to the IR.

The fine-registration process was only partly successful. This was quite apparent when viewing the output sequences in a movie format.   One observes in the movie sequence a considerable amount of jitter due to the random misalignment of the visible image. These effects are unlikely to appear when the two sensors have the same field of view (as opposed to the large difference that existed in the set of video segment used in this study).

3.0 IMAGE DECOMPOSITION FOR MULTI-SCALE FUSION

We have implemented a multi scale image fusion algorithm. Each of the input images is divided into spatial sub bands of different scales. Within each scale a combination rule for the corresponding pixels taken from the multiple (one per sensor input) components is applied. However, we use a contrast measure to determine the combination rule and apply the rule to the amplitude image (in each scale). These changes offer higher sensitivity to visual image aspects and assure reconstruction of a single channel if the other channel has no signal or very weak signal. In addition, this approach permits contribution for each pixels to vary among the source images depending on scale, thus even when a binary combination rule is applied, each pixel of the combined fused image may have a contributions to its values from all sensors.

We have compared the use of isotropic and oriented multi-scale oriented filters in the multi-scale decomposition. We first calculated a set of isotropic filters to obtain band-pass amplitude scale representations. This kind of decomposition is similar to the image decomposition that takes place in the eye's retina. We also applied four oriented filters that result in a set of scaled oriented images that are motivated by current cortical visual system models. In addition, we have tested the use of two oriented filters. The filters are designed to permit complete image reconstruction by simple sum of all the filtered oriented versions. Normalization of the amplitude signals by the local luminance completed the decomposition into visually relevant scale representations –local band limited contrast (Peli 1990).

3.1 Filter Implementation

Isotropic one- octave band-pass filters have been implemented along with a DC low-pass band and a high-pass band. One-dimensional representation of these filters for a is shown in Figure 3.1 on both a linear and log (base 2) frequency scale. For perfect reconstruction of the input images, filters sum to unity. Thus, a summation of all of the scales from single image decomposition reconstructs the original image. The equation used to generate the filters for the r^{th} spatial frequency at the i^{th} scale is:

$$G_i(r) = 0.5\left[1 + \cos\left(\pi \log_2 r - \pi i\right)\right]$$

To implement the filters in a two dimensional plane, r must now represent the distance from the origin of the Fourier domain, or at pixel $(f_x f_y)$. Figure 3.2 depicts filters for a choice of four orientations.
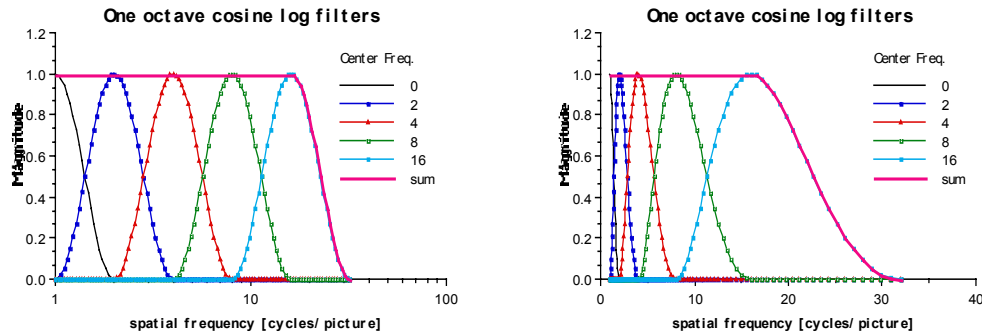

Figure 3.1: Isotropic filters response functions on linear (left) and log (right) scale
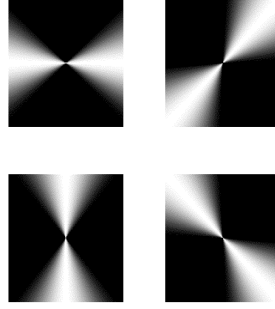
Figure 3.2: Oriented filters for four orientations

3.2   Contrast-based Image Fusion

The contrast based fusion algorithm calculates, for each of the oriented bandpass filtered versions of the images, a corresponding oriented band-pass contrast image similar to the isotropic contrast measure computed by Peli (1990). For each pixel at each scale i, for each spectral band (electromagnetic) j, and filter orientation l, the contrast measure is,

$$c_{i,j,l}(x, y) = \frac{a_{i,j,l}(x, y)}{l_i(x, y)}$$

where $a_{i,j,l}(x,y)$ is the band-pass filtered image in the i^{th} spatial frequency octave, the j^{th} orientation and the l^{th} band and $l_{i,l}(x,y)$ is the lowpass filtered image which represents the local luminance mean.

The basic fusion process compares the calculated contrast measure for each pixel, at each scale and orientation, and selects the spectral band that should dominate these variables in the fused image. The simplest binary rule will select the component with the higher contrast.   Once a contrast component $c_{i,j,l}(x,y)$ is selected the corresponding amplitude component $a_{i,j,l}(x,y)$ is added to the fused image. As can be seen if one spectral component has no signal, and thus no contrast 9or just very low contrast, this process will reconstruct the active spectral band image.

A number of variations of the basic algorithm were evaluated. The use of 4 orientations was found to provide the best performance (least artifacts) versus two orientations (horizontal and vertical) or isotropic filtering alone. It is noted that the gain from two orientations to four is not overwhelming. As a result, in analyzing the requirements for a real-time implementation, it is likely that the saving in processing only two orientations far outweighs any performance gain.
We implemented a strong preference for the intensified visible band at low scales representing large features. This was done under the assumption that large features are likely to be surface landmarks needed for pilotage and navigation. Thus, their correct interpretation of shape from is needed mostly for piloting tasks shading under stressful situations.

4.0 RESULTS - MONOCHROME FUSION

Figures 4.1 through 4.3 show the images intensified (II), IR and fused images for a sequence denoted by "up3".

Here, we can see that the object, marked in the output image, is transferred from the I^2 image. Also, information from both images is used to provide better detail to the road and trees. The fused image maintains the shape from shading of the visual band for large piloting related objects (hills and valleys). This preference for the visual band does not preclude the use of high frequency small features from the IR.

Figure 4.1: up3 I^2

Figure 4.2: up3 IR

Figure 4.3: up3 Fused Image

Figures 4.1 through 4.3 show the images intensified (II), IR and fused images for the up3 image. The bright object, marked in the output image, is transferred from the II image, while the dark clouds next to it are from the IR image. Information from both images is used to provide better detail to the road and trees. The 3-D shape-from-shading information available in the II image and not in the IR image is preserved in the fused image.

5.0 POLARITY PRESERVING COLOR CODING

The polarity preserving color coding implemented was aimed at providing the pilot with source information regarding fused image features by assigning color only to pixels where the luminance was controlled by the IR image. In the basic implementation the color assignment was applied while maintaining the luminance of the monochrome fused image. The color coding thus only serves to indicate which features in the fused image were

derived from the IR sensor, and what was their polarity relative to their surrounding in that image. The latter information should aid the interpretation by marking the object as hotter or colder than its surrounding background.

In most display systems a signal with equal values of R, G, and B inputs results in gray pixel. However, since the luminance component of an RGB signal is calculated as [Peli 1992]

$$Y= 0.299R + 0.587G + 0.114B$$

these proportions will have to be maintained to preserve the luminance in the final color presentation of the fused image. For each luminance value Y, we will compute the maximal red representation of the same pixel without changing the luminance, and without requiring negative inputs in either of the other color channels.

For a given gray scale value x out of (for example) 256, this requires that the factor z, by which the green and blue components should be modified for a positive red contrast pixel, satisfies the equation

$$z= (x - 256 * 0.299) / 0.701 * x$$

z will be positive for values of x above 77. Thus for pixels of gray scale above 77, the red channel will be assigned the highest value. However, for pixels of lower level luminance, the red component will be proportionally reduced from the maximal 256 to create the complete lookup table. The calculation for the cyan pixels representing negative contrast in the IR band is similar. To preserve luminance, red coloring ranges from black to red to orange to yellow to white. Cyan ranges from black to cyan to white.

Another possible approach is to relax the luminance preservation requirement and force the coloring of positive contrast pixels to vary from black to red. It is also more intuitive that for the most negative contrast (the zero value - coldest) to be most cyan, and as level increases move towards gray. In both the red and cyan case here, the goal is to not have neutral levels (128 out of 255) stand out in the fused image.

The assignment of color was addressed with an approach (geometry-based) that requires a parallel process to generate a binary map of target sized objects. The morphology-based approach utilizes Atlanticís target detection process [Peli 1993] that has demonstrated robust performance across multiple applications and sensor domains and requires only the specification of target size.

6.0 SUMMARY

A robust multi-spectral fusion process that consists of two stages was developed. The monochrome fusion stage was shown to capture the details from all of its inputs across a diverse set of inputs using a single algorithm setting. The algorithm maintains the natural visual appearance of large scale terrain features which are necessary for navigation and pilotage and increase the visibility of small scale targets obtained from both bands. A color coding, which followed this fusion serve to highlight target and at the same time indicate to the pilot the source of the target and its polarity in the more sensitive sensor.

ACKNOWLEDGEMENTS

REFERENCES

[AAEC 1997] "Practical AGC for Quasi-Invariant Target Signature," Fin. Rep. to Army/MICOM.
[Braje 1997] Braje W.L. " Do shadows influence recognition of natural objects" (ARVO 1997) Investigative Ophthalmol. &Vis. Sci. Suppl. pg s1002.

[Georgeson 1992] Georgeson M.A. "Human vision combines oriented filters to compute edges" Proc. R. Soc. Lond. B 235-245 1992.

[Gur 1993] Gur M and Syrkin G Color enhances Mach Band s detection threshold and percieved brightness Vision Res 33; 2313-2319 1993

[Jenkin 1997]   Jenkin and Howard "Interpreting shape from shading under different visual frames and body orientation" (ARVO 1997) Investigative Ophthalmol. &Vis. Sci. Suppl. pg s1002.

[Katz 1987]   Katz et al, "Application of Spectral Filtering to Missile Detection Using Staring Sensors at MWIR Wavelengths", Proceedings of the IRIS Conf. on Targets, Backgrounds, and Discrimination, Feb. 1987

[Morrone 1989], Morrone M.C. and D.C.,"Discrimination of spatial phase in central and peripheral vision," *Vision Research* 29 : 433-445.

[Morgan 1991] Morgan M.J., Ross J., and Hayes A. "The relative importance of local phase and local amplitude in patchwise image reconstraction" Biol. Cybern. 63:113-119, 1991.

[Paval 1991]   Paval M., Larimer J., and Ahumanda A. " Sensor fusion for synthetic vision" AIAA conference "Computing in Aerospace 8" Baltimore, MD, October 21-24, 1991.

[Peli 1990]   Peli E "Contrast in complex images," J. Opt. Soc. Am. A, 7: 2032-2040, 1990

[Peli 1992] Peli E, "Display nonlinearity in digital image processing," Opt. Engineering 31; 1992.

[Peli 1997c] Peli E (1997) In search of a contrast metric: matching the percieved contrast of Gabor patches at different phases and bandwidths. Vision Research In press.

[Peli 1979]   Peli, T., and Lim, J. S., "Adaptive Filtering For Image Enhancement," Proc. of ICASSP Conf., 1979.

[Peli 1993]   Peli, T., Vincent L.,Tom., V., *Morphology-based Detection and Segmentation in FLIR Imagery*, SPIE Conf. on Architecture, Hardware, and FLIR Issues in ATR, Orlando, FL, 1993.

[Peli 1997d] Peli, T., Monsen, P., Stahl, R., and Pauli, M., Long Range Detection and Tracking of Missiles, IRIS/IRCM, May 1997

[Peli 1997a]   Peli, T.,Young, M., and Ellis, K" Combining Linear and Nonlinear Processes for Multi-Spectral Material Detection/Identification," SPIE Conf. on Algorithms for Multispectral and Hyperspectral Imagery II, 1997 .

[Peli 1997b] Peli, T.,Young, M., and Ellis, K.,"Multi-Spectral Change Detection," SPIE Conf. on Algorithms for Multispectral and Hyperspectral Imagery II, Orlando, FL, 1997.

[Perconti] Perconti P., and Steele P.M. "Part Task Investigation of Multispectral Image Fusion using Gray Scale and Synthetic Color Night Vision Sensor Imagery for Helicopter Pilotage", undated report.

[Ryan 1995] Ryan D., and Tinkler R. " Night pilotage assessment of image fusion" SPIE vol 2465, 1995.

[Scribner 1993] Scribner D. A., Satyshur M.F., and Kruer M.R. "Composite Infrared color images and related processing" IRIS Targets, Back. and Discrim. January 1993, San Antonio TX.

[Scrofani 1997] Scrofani J.W. "An adaptive method for the enhanced fusion of low-light visible and uncooled thermal infrared imagery" MSc thesis 1997

[Stotts 1990]   Stotts, L.B., Winter, E.M. and Reed, I.S., "Clutter Rejection Using Multispectral Processing," SPIE Aerospace Sensing, Conference 1305 (16-20 April) 1990, Orlando, FL.

[Subramaniam 1997] Subramaniam and Biederman "Effect of Contrast Reversal on Object Recognition (ARVO 1997) Investigative Ophthalmol. & Vis. Sci. Suppl.

[Toet 1992]   Toet A " Multiscale contrast enhancement with application to image fusion" Optical Engineering 1992 vol 31 1026-1031.

[Waxman 1995] Waxman M.A., Fay D.A., Gove A.N., Seibert M., Racamato J.P., Carrick J.E., and SavoyeE.D. "Color Night Vision: Fusion of Intensified Visible and Thermal IR Imagery" SPIE Vol 2463, 58-68 1995.